arXiv:1703.06324v1 [cs.IR] 18 Mar 2017

Deep Tensor Encoding

B Sengupta*
Dept. of Engineering, University of Cambridge
Trumpington Street, Cambridge CB2 1PZ, UK
bs573@cam.ac.uk

E Vasquez and Y Qian†
Cortexica Vision Systems Limited
Capital Tower – 91 Waterloo Road, London SE1 8RT, UK
eduard.vazquez,yu.qian@cortexica.com

March 21, 2017

Abstract

Learning an encoding of feature vectors in terms of an over-complete dictionary or a probabilistic information geometric (Fisher vectors) construct is wide-spread in statistical signal processing and computer vision. In content based information retrieval using deep-learning classifiers, such encodings are learnt on the flattened last layer, without adherence to the multi-linear structure of the underlying feature tensor. We illustrate a variety of feature encodings incl. sparse dictionary coding and Fisher vectors along with proposing that a structured tensor factorization scheme enables us to perform retrieval that is at par, in terms of average precision, with Fisher vector encoded image signatures. In short, we illustrate how structural constraints increase retrieval fidelity.

1 Introduction

The success of deep-learning lies in constructing feature spaces where in competing classes of objects, sounds, etc. can be shattered using high-dimensional hyperplanes. The classifier relies on the accumulation of representation in the final convolution layer of a deep neural network. Often times, the classifier performance increases as one incorporates information from earlier layers of the neural network. Such a structural constraint has been imposed on certain deep learning architectures via the *inception module*. In addition to decreasing the computational effort, utilization of 1×1 convolution filters enables the dimensionality of feature map to be immensely reduced; in tandem with pooling, the dimensionality reduces even further. Thus, information from the previous layer(s) can be accumulated and concatenated in the inception module. By learning the weights feeding to the inception module there is the additional flexibility in vetting the different sources of information that reaches deeper layers.

A big demerit of deep-learning architectures is their inability to perform well in the advent of small amounts of training data. Tricks such as input (rotation, blur, etc.) and feature augmentation (in terms of inception module) have proven useful [8]. Such structural constraints regularize the optimization problem, reducing over-fitting. In this paper, we propose a much simpler structural constraint i.e., to utilize the multi-linear structure of deep feature tensors. We will first emphasize the importance of feature encoding, starting with Fisher encoding (of the last layer) followed by a sparse dictionary based on the last feature layer; this feeds to deep tensor factorization that brings together tensor learning and deep learning – ce-

*BS has dual appointments at Cortexica Vision Systems Ltd. and the Dept. of Bioengineering, Imperial College London.

†YQ and BS contributed equally to this contribution.

menting the idea that taking into account the high dimensional multi-linear representation increases the fidelity of our learnt representation. The algorithms are evaluated on a content-based-image-retrieval (CBIR) setting.

2 Methods

2.1 Dataset and deep-feature generation

In this paper, Describable Textures Dataset (DTD) [4] is used to evaluate feature encodings for image retrieval. Rather than recognition and description of object, texture images in DTD are collected from wild images (Google and Flickr) and classified based on human visual perception [16], such as directionality (line-like), regularity (polka-dotted and chequered), etc. DTD contains 5640 wild texture images with 47 describable attributes drawn from the psychology literature, and is publicly available on the web at http://www.robots.ox.ac.uk/vgg/data/dtd/.

Textures can be described via orderless pooling of filter bank response [7]. In Deep convolutional neural network (CNN), the convolutional layers are akin to non-linear filter banks; these have in fact been proved to be better for texture descriptions [5]. Here, the local deep features are extracted from last convolutional layer of a pre-trained VGG-M [3]. This is represented by $T = \left(t_1, ...t_i, ..., t_N : t \in \mathbb{R}^D\right)$; the size of last convolutional layer is $H \times W \times D$, where D denotes the dimension of filter response at the i^{th} pixel of last convolution layer; $N = H \times W$ is the total number of local features. Different feature encoding strategies are then utilized for encoding local descriptors. A similarity metric is then applied to rank images. We use the l_2 norm (Frobenius norm for matrices and tensors) between vectors as a notion of distance between the query and the database images.

We will now introduce five different encodings of the feature matrix – (a) Fisher encoding, (b) Sparse matrix dictionary learning, (c) t-SVD factorization, (d) Low-rank plus Sparse factorization, and e) Multilinear Principal Component Analysis (MPCA).

Feature encoding

2.1.1 Fisher encoding

We use a Gaussian Mixture Model (GMM) for encoding a probabilistic visual vocabulary on the training dataset. Images are then represented as Fisher Vectors [14, 17] – derivatives of log-likelihood of the model with respect to its parameters (the score function). Fisher encoding describes how the distribution of features of an individual image differs from the distribution fitted to the feature of all training images.

First, a set of D dimension local deep features are extracted from an image and denoted as $T = \left(t_1, ...t_i, ..., t_N : t \in \mathbb{R}^D\right)$. As Ref. [15, 5], a K component GMM with diagonal covariance is generated on the training set with the parameters $\{\Theta = (\omega_k, \mu_k, \Sigma_k)\}_{k=1}^K$, only the derivatives with respect to the mean $\{\mu_k\}_{k=1}^K$ and variances $\{\Sigma_k\}_{k=1}^K$ are encoded and concatenated to represent an image $T(X, \Theta) = \left(\frac{\partial L}{\partial \mu_1}, ..., \frac{\partial L}{\partial \mu_K}, \frac{\partial L}{\partial \Sigma_1}, ..., \frac{\partial L}{\partial \Sigma_K}\right)$, where,

$$\begin{aligned} L(\Theta) &= \sum_{i=1}^{N} log\left(\pi\left(t_i\right)\right) \\ \pi\left(t_i\right) &= \sum_{k=1}^{K} \omega_k N\left(t_i; \mu_k, \Sigma_k\right) \end{aligned} \tag{1}$$

For each component k, mean and covariance deviation on each vector dimension $j = 1, 2...D$ are

$$\begin{aligned} \frac{\partial L}{\partial \mu_{jk}} &= \frac{1}{N\sqrt{\omega_k}} \sum_{i=1}^{N} q_{ik} \frac{t_{ji} - \mu_{jk}}{\sigma_{jk}} \\ \frac{\partial L}{\partial \Sigma_{jk}} &= \frac{1}{N\sqrt{2\omega_k}} \sum_{i=1}^{N} q_{ik} \left[\left(\frac{t_{ji} - \mu_{jk}}{\sigma_{jk}}\right)^2 - 1\right] \end{aligned} \tag{2}$$

where q_{ik} is the soft assignment weight of feature t_i to the k^{th} Gaussian and defined as

$$q_{ik} = \frac{exp\left[-\frac{1}{2}\left(t_i - \mu_k\right)^T \Sigma_k^{-1}\left(t_i - \mu_k\right)\right]}{\sum_{t=1}^{K} exp\left[-\frac{1}{2}\left(t_i - \mu_t\right)^T \Sigma_t^{-1}\left(t_i - \mu_t\right)\right]} \tag{3}$$

Just as the image representation, the dimension of Fisher vector is $2KD$, K is the number of components

in the GMM, and D is the dimension of local feature descriptor. After l_2 normalization on Fisher vector, the Euclidean distance is calculated to find similarity between two images.

2.1.2 Sparse coding on deep features

The compressed (sparse) sensing framework allows us to learn a set of over-complete bases D and sparse weights ϕ such that the feature matrix T can be represented by a linear combination of these basis vectors:

$$\arg\min_{D,\phi} \frac{1}{n} \sum_{i=1}^{n} \left(\frac{1}{2} \left\| T - D \cdot \phi_i \right\|_F^2 \ + \lambda \left\| \phi_i \right\|_1 \right) \tag{4}$$

The k-SVD algorithm [9] comprises of two stages – first, a sparse coding stage (either using matching pursuit or basis pursuit) and second, a dictionary update stage. In the first stage when the dictionary D is fixed, we solve for ϕ using orthogonal matching pursuit (OMP). Briefly, OMP recovers the support of the weights ϕ using an iterative greedy algorithm that selects at each step the column of D that is most correlated with the current residual. Practically, we initialise the residual r_k, subsequently computing the column that reduces the residual the most j^* and finally adding this column to the support I_k

$$\begin{aligned} r_k &= T - D\phi_{k-1} \\ j^* &= \arg\min_{\substack{j=1\ldots n \\ \phi}} \left\| r_k - d_j \phi \right\|_2 \\ I_k &= I_{k-1} \cup j^* \end{aligned} \tag{5}$$

Iterating through these equations for a pre-specified number of iteration, we can update the sparse weight ϕ. After obtaining the sparse weights, we use a dictionary update stage where we update only one column of D each time. The update for the k-th column is,

$$\begin{aligned} \left\| T - D \cdot \phi \right\|_F^2 &= \left\| T - \sum_{j=1}^{K} D_j \cdot \phi_j^T \right\|_F^2 \\ &= \left\| \underbrace{\left(T - \sum_{j \neq k} D_j \cdot \phi_j^T \right)}_{E_k} - D_k \cdot \phi_k^T \right\|_F^2 \end{aligned} \tag{6}$$

In order to minimize $\left\| E_k - D_k \cdot \phi_k^T \right\|_F^2$ we decompose (SVD) E_k as UWV^T. Using this decomposition we utilize a rank-1 approximation to form $d_k = u_0$ and $\phi_k = w_0 v_0$. We can then iterate this for every column of D. Sparsity is attained by collapsing each ϕ_k only to non-zero entries, and constraining E_k to the corresponding columns.

For image retrieval, each local deep feature can be encoded by sparse weights ϕ. Such an image signature can be represented by max pooling of a set of ϕ, followed by measuring a distance between such sets.

2.1.3 Tensor factorization of deep features

In the earlier section, we relied on an alternate minimization of the dictionary and the loadings (weights) to yield a convex optimization problem. Specifically, the last convolutional layer was used to form a dictionary without reference to the tensorial (multi-linear) representation of feature spaces obtained from the deep convolutional network. Thus, our goal is to approximate these tensors as a sparse conic combinations of atoms that have been learnt from a dictionary comprising the entire image training set. In other words, we would like to obtain an over-complete dictionary such that each feature tensor can be represented as a weighted sum of a small subset of the atoms (loadings) of the dictionary.

There are two decompositions that are used for factorizing tensors, one is based on Tucker decomposition whilst the second is known as Canonical Decomposition/Canonical Polyadic Decomposition (CANDECOMP/CPD), also known as Parallel Factor Analysis (PARAFAC). In the former, tensors are represented by sparse core tensors with block structures. Specifically,

$\mathcal{T}$ is approximated as a multi-linear transformation of a small "core" tensor $\mathcal{G}$ by factor matrices A [11],

$$\mathcal{T} = \mathcal{G} \bullet_1 A^{(1)} \bullet_2 A^{(2)} \ldots_N A^{(N)} \triangleq \left[\left[\mathcal{G}; A^{(1)}, A^{(1)} \ldots A^{(N)}\right]\right] \quad (7)$$

In the latter, a tensor $\mathcal{T}$ is written as a sum of R rank-1 tensors, each of which can be written as the outer product of N factor variables i.e.,

$$\mathcal{T} = \sum_{r=1}^{R} a_r^{(1)} \otimes a_r^{(2)} \ldots a_r^{(N)} \triangleq \left[\left[A^{(1)}, A^{(2)} \ldots A^{(N)}\right]\right] \quad (8)$$

It is canonical when the rank R is minimal; such a decomposition is unique under certain conditions [6]. Even then, due to numerical errors, factorization of a feature matrix obtained using a deep neural network results in a non-unique factorization. Thus, CPD proves inadequate for unique feature encoding. We therefore utilize a factorization that is similar to a 2D-PCA albeit lifted for multilinear objects. Specifically, we use t-SVD [10] to factorize the feature matrices.

Based on t-SVD

The t-product between two tensors, $\mathcal{T}_1$ and $\mathcal{T}_2$,

$$\begin{aligned} \mathcal{T}_1 * \mathcal{T}_2 &\equiv \text{fold}(\text{circ}(\mathcal{T}_1) \cdot \text{unfold}(\mathcal{T}_2)) \\ \mathcal{T} &= \mathcal{U} * \mathcal{S} * \mathcal{V}^T \end{aligned} \quad (9)$$

where, circ$(\cdot)$ creates a block circulant matrix and the unfold operator matricizes the tensor on the tube (3rd) axis. $\mathcal{S}$ is a f-diagonal tensor that contains the eigentupules of the covariance tensor on its diagonal whereas, the columns of $\mathcal{U}$ are the eigenmatrices of the covariance tensor. The images in the training set are inserted as the second index of the tensor. In other words, using t-SVD, we obtain an orthogonal factor dictionary of the entire training set. Ultimately, a projection of the mean-removed input tensor (i.e., a single feature tensor, $\mathcal{T}_{\text{test}}$) on the orthogonal projector ($U^T * \mathcal{T}_{\text{test}}$) forms the tensor encoding of each image. The Frobenius norm then measures the closeness between a pair of images (individual tensor projections). Computation efficiency is guaranteed since the comparison between the test image and the database is measured in Fourier domain – the t-product utilizes a fast Fourier transform algorithm in its core [13].

Another feature encoding that we consider is the partition of each individual tensor i.e., $\mathcal{T} = \mathcal{L} + \mathcal{P}$ where, $\mathcal{L}$ is a low-rank tensor and $\mathcal{P}$ is a sparse tensor. We have $\mathcal{L} = \mathcal{U}_{1:r} * \mathcal{S}_{1:r} * \mathcal{V}_{1:r}^T$ and $\mathcal{P} = \mathcal{T} - \mathcal{L}$. r denotes the truncation index of the singular components.

Based on mPCA

For high-order tensor analysis, multilinear Principal Component Analysis(mPCA) or High Order Singular Value Decomposition(HOSVD) [18, 12] compute a set of orthonormal matrices associated with each mode of a tensor – this is analogous to the orthonormal row and column space of a matrix computed by the matrix PCA/SVD.

In a Tucker decomposition (Eqn. 7) if the factor matrices are constrained to be orthogonal the input tensor can be decomposed as a linear combination of rank-1 tensors. Given a set of N-order training tensor $\mathcal{T}$, the objective of mPCA is to find N linear projection matrices that maximize the total tensor scatter (variance) in the projection subspace. When factor matrices $A^{(N)}$ in Eqn. 7 are orthogonal then $\|\mathcal{T}\|_F^2 = \|\mathcal{G}\|_F^2$. Each query tensor $\mathcal{T}_q$ can then be projected using $\mathcal{Y} = \mathcal{T}_q \times_1 \mathbf{A}^{(1)T} \times_2 \mathbf{A}^{(2)T} \ldots \times_N \mathbf{A}^{(N)T}$ where the bold-faced matrices represent a low-dimensional space. The objective then becomes learning a set of matrices $\mathbf{A}$ to solve

$$\underset{\mathbf{A}^{(1)} \ldots \mathbf{A}^{(N)}}{\arg\max} \sum_{m=1}^{M} \left\| \mathcal{Y} - \overline{\mathcal{Y}} \right\|_F^2 \quad (10)$$

$\overline{\mathcal{Y}}$ is the mean tensor. Since the projection to an N^{th} order tensor subspace consists of N projections to N vector subspaces, in Ref. [12], optimization of the N projection matrices can be solved by finding each $A^{(n)}$ that maximizes the scatter in the n-mode vector subspace with all other A-matrices fixed. The local optimization procedure can be iterated until convergence.

Experiments

In this section, five deep feature encoding methods are evaluated for image retrieval on the DTD dataset. Images are resized to same size (256 x 256), deep feature

Figure 1: Retrieved results on DTD: a) Fisher vector b) Sparse coding c) t-SVD d) mPCA e) Low rank tensor f) Raw tensor

is extracted from last convolutional layer of a pre-trained VGG-M. For Fisher vector and sparse coding, deep features are flattened as a set of 1D local features. For t-SVD, deep features are represented as 2D feature maps, and treated as [HxW,1,D] data structures (see Methods). After encoding and l_2 normalization, the Euclidean distance is calculated to find similarity between two images.

To evaluate image retrieval, Mean Average Precision (MAP) on top 10 rankings are calculated. Two images per category i.e., a total of 94 images are selected as queries from the test dataset. The dataset retrieved includes 3760 images from DTD training and validation datasets. MAP on DTD is listed in Table 1. An example of the retrieval obtained with each method is shown in Figure 2. On each case 10 images are displayed. Top left image is the query used. The rest of images are arranged by similarity to query image as obtained with each encoding method.

Image retrieval amounts to searching large databases to return ranked list of images that have similar properties to that of the query image. Here, we reason that in contrast to raw feature tensors, their encodings either using Fisher vector, sparse dictionaries or multi-linear tensors increases the average precision of retrieval. Table 1 shows that multi-linear encodings based on t-SVD, mPCA or low-rank decomposition of individual images all have similar fidelity whilst performing very close to information geometric encoding of the feature vectors.

Sparse coding supersedes other methods in terms of average precision. This is because the dictionary learnt using k-SVD is matched to the underlying image distribution (i.e., by learning the sparse weights) whereas the tensor dictionaries (t-SVD or mPCA) use orthogonal matrices as dictionaries without the added constraint to finesse the weights, or to add additional structure in terms or sparsity, low-rank or non-negativity.

Table 1: Average precision for the DTD dataset.

Methods	**top-1**	**top-5**	**top-10**	**time taken**
Fisher Vector	0.56	0.52	0.48	12ms
Sparse Coding	0.62	0.49	0.44	35ms
t-SVD dictionary	0.53	0.42	0.38	2188ms
mPCA dictionary	0.53	0.43	0.39	5ms
Low Rank	0.51	0.42	0.38	967ms
Raw Tensor	0.41	0.35	0.31	

As shown in Table 1, computing mPCA tensor encodings are computationally efficient in contrast to sparse dictionaries or learning a probabilistic model for Fisher vector encodings. Combined with reasonable retrieval performance, tensor encodings of deep neural features make them a contender for commercial infrastructures.

The code was implemented in Matlab 2015a under linux with Intel Xeon CPU E5-2640 @ 2.00GHz and 125G RAM. Sandia's Tensor toolbox [1], KU Leuven's Tensorlab [19] and TFOCS [2] were used to implement our algorithms.

Conclusion

Feature encoding is crucial for comparing images (or videos, text, etc.) that are similar to one another. Our experiments show that whilst sparse encoding of a feature tensor proves to be the most efficient encoding for retrieval, having no encoding grossly reduces the average precision. Taking the multi-linear properties of the feature tensor improves retrieval, and the fidelity is at par with Fisher encoding. We further show that computing such multi-linear representation of the feature tensor is computationally much efficient than constructing a sparse dictionary or learning a probabilistic model.

The sparse dictionary encoding has the highest average precision due to the added flexibility of learning the weights as well as imposing the structural constraint of sparsity. Fisher vector encoding has the second highest precision because of its ability to capture the information

geometry of the underlying probabilistic model. Specifically, the Fisher tensor encodes the underlying Riemannian metric which augments the encoding with the curvature of the underlying distribution. The multi-linear approaches based on t-SVD, mPCA and low-rank decomposition perform at par with Fisher vectors as they encode the third and higher order interaction in the feature tensors. Comparison of the compute time tells us that amongst all of the methods, encoding deep feature tensors using mPCA is the most time-efficient algorithm.

Our results have not exploited the geometry exhibited by the tensors, for example, one can calculate lagged covariance tensors from the feature tensor – these tensors exhibit a Riemann geometry due to their positive definite structure. Therefore building a dictionary of co-variance tensors using t-SVD, wherein an Augmented Lagrangian alternating direction method can be employed to learn a sparse representation, is the next viable step to our work. We anticipate that such a multi-linear overcomplete dictionary should at the very least have increased precision to that of the Fisher encoding method. In so far, the last convolutional layer was used to form a dictionary without reference to the earlier layers of the deep neural network. In fact a step forward would be to constrain the construction of an image signature with tensorial information from an earlier layer. The tensor methods rely on factorizing large tensors, especially those that emerge from deep neural networks. Yet, GPU implementation of the Fourier transform in the form of *cuFFT* enables us to build a scalable commercial solution.

References

[1] B. W. Bader, T. G. Kolda, et al. Matlab tensor toolbox version 2.6. Available online, February 2015. 5

[2] S. R. Becker, E. J. Candès, and M. C. Grant. Templates for convex cone problems with applications to sparse signal recovery. *Mathematical programming computation*, 3(3):165, 2011. 5

[3] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference*, 2014. 2

[4] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi. Describing textures in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014. 2

[5] M. Cimpoi, S. Maji, and A. Vedaldi. Deep filter banks for texture recognition and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2

[6] I. Domanov and L. D. Lathauwer. On the uniqueness of the canonical polyadic decomposition of third-order tensors - part II: uniqueness of the overall decomposition. *SIAM J. Matrix Analysis Applications*, 34(3):876–903, 2013. 4

[7] Y. Gong, L. Wang, R. Guo, and S. Lazebnik. Multi-scale orderless pooling of deep convolutional activation features. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VII*, pages 392–407, 2014. 2

[8] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. 1

[9] Z. Jiang, Z. Lin, and L. S. Davis. Learning a discriminative dictionary for sparse coding via label consistent K-SVD. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1697 – 1704, 2011. 3

[10] M. E. Kilmer, K. Braman, N. Hao, and R. C. Hoover. Third-order tensors as operators on matrices: A theoretical and computational framework with applications in imaging. *SIAM Journal on Matrix Analysis and Applications*, 34(1):148–172, 2013. 4

[11] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Rev.*, 51(3):455–500, 2009. 4

[12] H. Lu, K. N. Plataniotis, and A. N. Venetsanopoulos. mPCA: Multilinear principal component analysis of tensor objects. *IEEE Trans. Neural Network*, 1(19):18–39, 2008. 4

[13] Nvidia. CUDA CUFFT library, 2007. 4

[14] F. Perronnin and C. Dance. Fisher kernels on visual vocabularies for image categorization. In *IEEE Conference on Computer Vision and Pattern Recognitio*, 2006. 2

[15] K. Simonyan, O. M. Parkhi, A. Vedaldi, and A. Zisserman. Fisher Vector Faces in the Wild. In *British Machine Vision Conference*, 2013. 2

[16] H. Tamura, S. Mori, and T. Yamawak. Textural features corresponding to visual perception. *IEEE Transactions on Systems, Man and Cybernetics*, page 460 473, 1978. 2

[17] T. Uricchio, M. Bertini, L. Seidenari, and A. D. Bimbo. Fisher encoded convolutional bag-of-windows for efficient image retrieval and social image tagging. In *ICCV Workshops*, pages 1020–1026, 2015. 2

[18] M. Vasilescu and D. Terzopoulos. Multilinear subspace analysis for image ensembles. In *IEEE Conference on*

Computer Vision and Pattern Recognition, pages 93–99, June 2003. 4

[19] N. Vervliet, O. Debals, L. Sorber, M. Van Barel, and L. De Lathauwer. Tensorlab 3.0, Mar. 2016. Available online. 5

A Tensor Norms

Let t_{ijk} be the elements of tensor $\mathcal{T}$, then the Frobenius norm is $\|\mathcal{T}\|_F = \|vec(\mathcal{T})\|_2 = \sqrt{\sum_i \sum_j \sum_k t_{ijk}^2}$. The nuclear (trace) norm is defined as $\|\mathcal{T}\|_* = trace(\sqrt{\mathcal{T}^T\mathcal{T}}) = \sum_{i=1}^{\min\{m,n\}} \sigma_i$. σ are the singular values of $\mathcal{T}$.